

PAVAN BALAJI

Mathematics and Computer Science (MCS) Division, Argonne National Laboratory (ANL)

Phone (O): 630.252.3017 **Fax:** 630.252.5986

Email: balaji@mcs.anl.gov **Home Page:** <http://www.mcs.anl.gov/~balaji>

PROFESSIONAL PREPARATION

| | |
|---|---------|
| Ph.D. , Computer Science and Engg., Ohio State University | 06/2006 |
| M.S. , Computer Science and Engg., Ohio State University | 06/2003 |
| B.Tech. , Computer Science and Engg., Indian Institute of Technology (IIT), Madras | 08/2001 |

APPOINTMENTS

Argonne National Laboratory, Argonne, IL

| | |
|------------------------------|---------------------|
| Assistant Computer Scientist | Aug '08 to present |
| Post-doctoral Researcher | Aug '06 to July '08 |

University of Chicago, Chicago, IL

| | |
|-----------------|--------------------|
| Research Fellow | Mar '07 to present |
|-----------------|--------------------|

RESEARCH INTERESTS

My research broadly falls into the general areas of parallel and distributed computing. Specifically, I focus on three areas: (i) network-based parallel and distributed computing including high-speed interconnects (InfiniBand, 10-Gigabit Ethernet, Myrinet) and protocols (iWARP, TCP/IP, UDP/IP); (ii) parallel programming models and other middleware (MPI, sockets, file-systems, web-based data-center middleware); (iii) job scheduling and resource management (QoS, parameter sweep applications).

AWARDS AND ACCOMPLISHMENTS

-
1. *Best Paper Award*. International Supercomputing Conference (ISC), Hamburg, Germany, 2009.
 2. *Best Paper Award*. IEEE International Conference on Cluster Computing (Cluster), Tsukuba, Japan, 2008.
 3. *Outstanding Paper Award*. European PVM/MPI Users' Group Conference (EuroPVM/MPI), Dublin, Ireland, 2008.
 4. *Outstanding Paper Award*. International Supercomputing Conference (ISC), Dresden, Germany, 2008.
 5. *Storage Challenge Award (Large Systems Category)*. IEEE/ACM International Conference for High-Performance Computing, Networking, Storage and Analysis (SC), Reno, Nevada, 2007.
 6. *Los Alamos Director's Technical Achievement*, Los Alamos National Laboratory, 2005.
 7. *Outstanding Research Award*, Computer Science and Engineering, Ohio State University, 2005.

RESEARCH GRANTS, DONATIONS AND LOANERS

-
1. **Co-PI** (PI: W. Gropp, Argonne National Laboratory): *Analyzing and Profiling Asymmetric MPI Collective Communication on BGW*. Supercomputing time on the 40,960 processor Blue Gene Watson (BGW) system. Period: 11/01/2007.
 2. **PI**: *Designing MPICH2 over Qlogic InfiniBand*. Equipment loaner from Qlogic Corporation. Period: 09/2007 to present. Total equipment cost: \$30,000.

3. **Co-PI** (PI: D. K. Panda, Ohio State University): *Designing Next Generation Communication and I/O Subsystems with Multicore Architectures*. National Science Foundation (NSF) Computing Processes and Artifacts (CPA) program. Period: 07/01/07 to 06/30/10. Total funding: \$462,512; Funding share: \$87,512.
4. **PI**: *10-Gigabit Ethernet Evaluation with MPI and PVFS*. Equipment loaner from Fulcrum Microsystems. Period: 12/2006 to present. Total equipment cost: \$15,000.
5. **PI**: *Designing Efficient System Software Services for iWARP-enabled Clusters*. Equipment loaner from NetEffect Inc. Period: 10/2006 to 02/2008. Total equipment cost: \$30,000.

SELECTED PUBLICATIONS

Book Chapters

1. D. K. Panda, **P. Balaji**, S. Sur and M. Koop. *Attaining High Performance Communication: A Vertical Approach*. Chapter on *Commodity High Performance Interconnects*; CRC Press, 2009.
2. W. Feng and **P. Balaji**. *Attaining High Performance Communication: A Vertical Approach*. Chapter on *Ethernet vs. Ethernet*; CRC Press, 2009.
3. **P. Balaji**, P. Sadayappan and M. Islam. *Market-Oriented Grid and Utility Computing*. Chapter on *Techniques on Providing Hard Quality of Service Guarantees in Job Scheduling*; Wiley Publishers, 2008.

Refereed Journal Articles

4. **P. Balaji**, A. Chan, R. Thakur, W. Gropp and E. Lusk. *Toward Message Passing for a Million Processes: Characterizing MPI on a Massive Scale Blue Gene/P*. Accepted for publication at the special edition of the Springer Journal of Computer Science on Research and Development (presented at the International Supercomputing Conference (ISC)), 2009. **Best Paper Award** at ISC.
5. P. Lai, **P. Balaji**, R. Thakur and D. K. Panda. *ProOnE: A General Purpose Protocol Onload Engine for Multi- and Many-Core Architectures*. Accepted for publication at the special edition of the Springer Journal of Computer Science on Research and Development (presented at the International Supercomputing Conference (ISC)), 2009.
6. **P. Balaji**, W. Feng and D. K. Panda. *Bridging the Ethernet-Ethernet Performance Gap*. IEEE Micro Journal; Special Issue on High-Performance Interconnects, pp. 24-40, Vol. 26, Issue 3, 2006.
7. H. -W. Jin, **P. Balaji**, C. Yoo, J. -Y. Choi and D. K. Panda. *Exploiting NIC Architectural Support for Enhancing IP based Protocols on High Performance Networks*. Special Issue of the Journal of Parallel and Distributed Computing (JPDC) on Design and Performance of Networks for Super-, Cluster- and Grid-Computing, pp. 1348-1365, Vol. 65, Issue 11, 2005.
8. M. Islam, **P. Balaji**, P. Sadayappan and D. K. Panda. *QoPS: A QoS based scheme for Parallel Job Scheduling (extended journal version)*. IEEE Springer LNCS Journal Series, pp. 252-268, Vol. 2862, 2003.

Refereed Conference Publications

9. G. Santhanaraman, **P. Balaji**, K. Gopalakrishnan, R. Thakur, W. Gropp and D. K. Panda. *Natively Supporting True One-sided Communication in MPI on Multi-core Systems with InfiniBand*. IEEE International Symposium on Cluster Computing and the Grid (CCGrid). May 18-21, 2009, Shanghai, China.
10. **P. Balaji**, S. Bhagvat, R. Thakur and D. K. Panda. *Sockets Direct Protocol for Hybrid Network Stacks: A Case Study with iWARP over 10G Ethernet*. International Conference on High Performance Computing (HiPC). Dec. 17-20, 2008, Bangalore, India.
11. A. Chan, **P. Balaji**, W. Gropp and R. Thakur. *Communication Analysis of Parallel 3D FFT for Flat Cartesian Meshes on Large Blue Gene Systems*. International Conference on High Performance Computing (HiPC). Dec. 17-20, 2008, Bangalore, India.

12. M. Kumar, V. Chaube, **P. Balaji**, W. Feng and H.-W. Jin. *Making a Case for Proactive Flow Control in Optical Circuit-Switched Networks*. International Conference on High Performance Computing (HiPC). Dec. 17-20, 2008, Bangalore, India.
13. H. Lin, **P. Balaji**, R. Poole, C. Sosa, X. Ma and W. Feng. *Massively Parallel Genomic Sequence Search on the Blue Gene/P Architecture*. IEEE/ACM International Conference for High Performance Computing, Networking, Storage and Analysis (SC). Nov. 15-21, 2008, Austin, Texas.
14. T. Scogland, G. Narayanaswamy, **P. Balaji** and W. Feng. *Asymmetric Interactions in Symmetric Multi-core Systems: Analysis, Enhancements and Evaluation*. IEEE/ACM International Conference for High Performance Computing, Networking, Storage and Analysis (SC). Nov. 15-21, 2008, Austin, Texas.
15. N. Desai, **P. Balaji**, P. Sadayappan and M. Islam. *Are Non-Blocking Networks Really Needed for High-End-Computing Workloads?* IEEE International Conference on Cluster Computing (Cluster), **Best Paper Award**, Sep. 29 - Oct. 1st, 2008, Tsukuba, Japan.
16. **P. Balaji**, A. Chan, W. Gropp, R. Thakur and E. Lusk. *Non-Data-Communication Overheads in MPI: Analysis on Blue Gene/P*. The Euro PVM/MPI Users' Group Conference (EuroPVM/MPI), **Outstanding Paper Award**, Sep. 7-10, 2008, Dublin, Ireland.
17. **P. Balaji**, D. Buntinas, D. Goodell, W. Gropp and R. Thakur. *Toward Efficient Support for Multithreaded MPI Communication*. The Euro PVM/MPI Users' Group Conference (EuroPVM/MPI), Sep. 7-10, 2008, Dublin, Ireland.
18. J. L. Traff, A. Ripke, C. Siebert, **P. Balaji**, R. Thakur and W. Gropp. *A Simple, Pipelined Algorithm for Large, Irregular All-gather Problems*. The Euro PVM/MPI Users' Group Conference (EuroPVM/MPI), Sep. 7-10, 2008, Dublin, Ireland.
19. G. Narayanaswamy, **P. Balaji** and W. Feng. *Impact of Network Sharing in Multi-core Architectures*. IEEE International Conference on Computer Communication and Networks (ICCCN), Aug. 3-7, 2008, St. Thomas, U.S. Virgin Islands.
20. **P. Balaji**, W. Feng and H. Lin. *Semantics-based Distributed I/O with the ParaMEDIC Framework*. ACM/IEEE International Symposium on High Performance Distributed Computing (HPDC), Jun. 23-27, 2008, Boston, Massachusetts.
21. **P. Balaji**, W. Feng, H. Lin, J. Archuleta, S. Matsuoka, A. Warren, J. Setubal, E. Lusk, R. Thakur, I. Foster, D. S. Katz, S. Jha, K. Shinpaugh, S. Coghlan and D. Reed. *Distributed I/O with ParaMEDIC: Experiences with a Worldwide Supercomputer*. International Supercomputing Conference (ISC). **Outstanding Paper Award**. Jun. 17-20, 2008, Dresden, Germany.
22. **P. Balaji**, W. Feng, J. Archuleta, H. Lin, R. Kettimuthu, R. Thakur and X. Ma. *Semantics-based Distributed I/O for mpiBLAST*. Short paper. ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP), Feb. 20-23, 2008, Salt Lake City, Utah.
23. **P. Balaji**, W. Feng, S. Bhagvat, D. K. Panda, R. Thakur and W. Gropp. *Analyzing the Impact of Supporting Out-of-Order Communication on In-order Performance with iWARP*. IEEE/ACM International Conference for High Performance Computing, Networking, Storage and Analysis (SC), Nov. 10-16, 2007, Reno, Nevada.
24. **P. Balaji**, W. Feng, J. Archuleta and H. Lin. *ParaMEDIC: Parallel Metadata Environment for Distributed I/O and Computing*. IEEE/ACM International Conference for High Performance Computing, Networking, Storage and Analysis (SC). **Storage Challenge Award**. Nov. 10-16, 2007, Reno, Nevada.
25. G. Narayanaswamy, **P. Balaji** and W. Feng. *An Analysis of 10-Gigabit Ethernet Protocol Stacks in Multicore Environments*. IEEE International Symposium on High-Performance Interconnects (HotI), Aug. 22-24, 2007, Palo Alto, California.
26. **P. Balaji**, S. Bhagvat, D. K. Panda, R. Thakur and W. Gropp. *Advanced Flow-control Mechanisms for the Sockets Direct Protocol over InfiniBand*. IEEE International Conference on Parallel Processing (ICPP), Sep. 10-14, 2007, Xi'an, China.
27. M. Islam, **P. Balaji**, G. Sabin and P. Sadayappan. *Analyzing and Minimizing the Impact of Opportunity Cost in QoS-aware Job Scheduling*. IEEE International Conference on Parallel Processing (ICPP), Sep. 10-14, 2007, Xi'an, China.

28. **P. Balaji**, D. Buntinas, S. Balay, B. Smith, R. Thakur and W. Gropp. *Nonuniformly Communicating Noncontiguous Data: A Case Study with PETSc and MPI*. IEEE International Parallel and Distributed Processing Symposium (IPDPS), Mar. 26-30, 2007, Long Beach, California.
29. **P. Balaji**, W. Feng, Q. Gao, R. Noronha, W. Yu and D.K. Panda. *Head-to-TOE Comparison for High Performance Sockets over Protocol Offload Engines*. IEEE International Conference on Cluster Computing (Cluster), Sep. 26-30, 2005, Boston, Massachusetts.
30. W. Feng, **P. Balaji**, C. Baron, L. N. Bhuyan and D. K. Panda. *Performance Characterization of a 10-Gigabit Ethernet TOE*. IEEE International Symposium on High Performance Interconnects (HotI), Aug. 17-19, 2005, Palo Alto, California.
31. S. Narravula, **P. Balaji**, K. Vaidyanathan, H. -W. Jin and D. K. Panda. *Architecture for Caching Responses with Multiple Dynamic Dependencies in Multi-Tier Data-Centers over InfiniBand*. IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGrid), May 9-12, 2005, Cardiff, UK
32. **P. Balaji**, K. Vaidyanathan, S. Narravula, H. -W. Jin and D. K. Panda. *On the Provision of Prioritization and Soft QoS in Dynamically Reconfigurable Shared Data-Centers over InfiniBand*. IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS), Mar. 20-22, 2005, Austin, Texas.
33. M. Islam, **P. Balaji**, P. Sadayappan and D. K. Panda. *Towards Provision of Quality of Service Guarantees in Job Scheduling*. IEEE International Conference on Cluster Computing (Cluster), Sep. 20-23, 2004, San Diego, California.
34. **P. Balaji**, S. Narravula, K. Vaidyanathan, S. Krishnamoorthy, J. Wu and D. K. Panda. *Sockets Direct Protocol over InfiniBand in Clusters: Is it Beneficial?*. IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS), Mar. 10-12, 2004, Austin, Texas.
35. R. Kurian, **P. Balaji** and P. Sadayappan. *Opportune Job Shredding: An Efficient Approach for Scheduling Parameter Sweep Applications*. Los Alamos Computer Science Institute (LACSI) Symposium, Oct. 12-14, 2003, Santa Fe, New Mexico.
36. **P. Balaji**, J. Wu, T. Kurc, U. Catalyurek, D. K. Panda and J. Saltz. *Impact of High Performance Sockets on Data Intensive Applications*. IEEE International Symposium on High Performance Distributed Computing (HPDC), Jun. 22-24, 2003, Seattle, Washington.
37. R. Gupta, **P. Balaji**, J. Nieplocha and D. K. Panda. *Efficient Collective Operations using Remote Memory Operations on VIA-based Clusters*. IEEE International Parallel and Distributed Processing Symposium (IPDPS), Apr. 22-26, 2003, Nice, France.
38. **P. Balaji**, P. Shivam, P. Wyckoff and D. K. Panda. *High Performance User-level Sockets over Gigabit Ethernet*. IEEE International Conference on Cluster Computing (Cluster), Sep. 23-26, 2002, Chicago, Illinois.

Refereed Workshop Publications

39. **P. Balaji**, S. Bhagvat, H. -W. Jin and D. K. Panda. *Asynchronous Zero-copy Communication for Synchronous Sockets in the Sockets Direct Protocol (SDP) over InfiniBand*. Workshop on Communication Architecture for Clusters (CAC); in conjunction with the IEEE International Parallel and Distributed Processing Symposium (IPDPS), Apr. 25th, 2006, Rhodes Island, Greece.
40. V. Viswanathan, **P. Balaji**, W. Feng, J. Leigh, D. K. Panda. *A Case for UDP Offload Engines in LambdaGrids*. Workshop on Protocols for Fast Long-Distance Networks (PFLDnet), Feb. 2-3, 2006, Nara, Japan.
41. **P. Balaji**, H. -W. Jin, K. Vaidyanathan and D. K. Panda. *Supporting iWARP Compatibility and Features for Regular Network Adapters*. Workshop on Remote Direct Memory Access (RDMA): Applications, Implementations and Techniques (RAIT); in conjunction with IEEE International conference on Cluster Computing (Cluster), Sep. 26th, 2005, Boston, Massachusetts.
42. H. -W. Jin, S. Narravula, G. Brown, K. Vaidyanathan, **P. Balaji** and D. K. Panda. *Performance Evaluation of RDMA over IP Networks: A Study with the Ammasso Gigabit Ethernet NIC*. Workshop on High Performance Interconnects for Distributed Computing (HPI-DC); in conjunction with IEEE International Symposium on High Performance Distributed Computing (HPDC), Jul. 24th, 2005, Research Triangle Park, North Carolina.

43. K. Vaidyanathan, **P. Balaji**, H. -W. Jin and D. K. Panda. *Workload driven analysis of File Systems in Shared Multi-Tier Data-Centers over InfiniBand*. Workshop on Computer Architecture Evaluation using Commercial Workloads (CAECW); in conjunction with IEEE International Symposium on High Performance Computer Architecture (HPCA), Feb. 12th, 2005, San Francisco, California.
44. **P. Balaji**, H. V. Shah and D. K. Panda. *Sockets vs. RDMA Interface over 10-Gigabit Networks: An In depth Analysis of the Memory Traffic Bottleneck*. Workshop on Remote Direct Memory Access (RDMA): Applications, Implementations and Technologies (RAIT); in conjunction with IEEE International Conference on Cluster Computing (Cluster), Sep. 20th, 2004, San Diego, California.
45. **P. Balaji**, K. Vaidyanathan, S. Narravula, S. Krishnamoorthy, H. -W. Jin and D. K. Panda. *Exploiting Remote Memory Operations to Design Efficient Reconfiguration for Shared Data-Centers over InfiniBand*. Workshop on Remote Direct Memory Access (RDMA): Applications, Implementations and Technologies (RAIT); in conjunction with IEEE International Conference on Cluster Computing (Cluster), Sep. 20th, 2004, San Diego, California.
46. S. Narravula, **P. Balaji**, K. Vaidyanathan, S. Krishnamoorthy, J. Wu and D. K. Panda. *Supporting Strong Coherency for Active Caches in Multi-Tier Data-Centers over InfiniBand*. Workshop on System Area Networks (SAN); in conjunction with IEEE International Symposium on High Performance Computer Architecture (HPCA), Feb. 14th, 2004, Madrid, Spain.
47. M. Islam, **P. Balaji**, P. Sadayappan and D. K. Panda. *QoPS: A QoS based scheme for Parallel Job Scheduling*. Job Scheduling Strategies for Parallel Processing (JSSPP) workshop; in conjunction with IEEE International Symposium on High Performance Distributed Computing (HPDC), Jun. 24th, 2003, Seattle, Washington.

Invited Papers

48. W. Feng, **P. Balaji** and A. Singh. *Network Interface Cards as First-Class Citizens*. Workshop on The Influence of I/O on Microprocessor Architecture (IOM); in conjunction with the IEEE International Symposium on High Performance Computer Architecture (HPCA), Feb. 15th, 2009, Raleigh, North Carolina.
49. K. Vaidyanathan, S. Narravula, **P. Balaji** and D. K. Panda. *Designing Efficient Systems Services and Primitives for Next-Generation Data-Centers*. Workshop on the National Science Foundation Next Generation Software (NSFNGS) Program; in conjunction with the IEEE International Parallel and Distributed Processing Symposium (IPDPS), Mar. 26th, 2007, Long Island, California.
50. **P. Balaji**, K. Vaidyanathan, S. Narravula, H. -W. Jin and D. K. Panda. *Designing Next Generation Data-centers with Advanced Communication Protocols and Systems Services*. Workshop on the National Science Foundation Next Generation Software (NSFNGS) Program; in conjunction with the IEEE International Parallel and Distributed Processing Symposium (IPDPS), Apr. 25th, 2006, Rhodes Island, Greece.

Technical Reports

51. K. Vaidyanathan, S. Bhagvat, **P. Balaji** and D. K. Panda. *Understanding the Significance of Network Performance in End Applications: A Case Study with EtherFabric and InfiniBand*. Technical Report, OSU-CISRC-2/06-TR19, Ohio State University, Feb, 2006.
52. K. Vaidyanathan, **P. Balaji**, J. Wu, H. -W. Jin and D. K. Panda. *An Architectural Study of Cluster-based Multi-tier Data-Centers*. Technical Report, OSU-CISRC-5/04-TR25, Ohio State University, May, 2004.
53. S. Krishnamoorthy, **P. Balaji**, K. Vaidyanathan, H. -W. Jin and D. K. Panda. *Dynamic Reconfigurability Support for providing Soft Quality of Service Guarantees in Multi-Tier Data-Centers over InfiniBand*. Technical Report, OSU-CISRC-2/04-TR10, Ohio State University, Feb, 2004.

Dissertation/Thesis

54. **P. Balaji**. *High Performance Communication Support for Sockets-based Applications over High-speed Networks*. Ph.D. Dissertation, Ohio State University, June, 2006.
55. **P. Balaji**. *Java Animated Text Rendering for Indian Language Scripts*. B.Tech. Thesis, Indian Institute of Technology, Madras, May, 2001.

PRESENTATIONS

Invited Talks

1. *And You Thought Threads Could Work in Parallel?* University of New Mexico (UNM), Oct 13th, 2008, Albuquerque, New Mexico.
2. *Interfacing Multicores and High-speed Networks Towards Large-scale Computing Systems.* Virginia Tech (VT), Dec 11th, 2007, Blacksburg, Virginia.
3. *Designing High-end Computing Systems with InfiniBand.* Illinois Institute of Technology (IIT), Jun 21st, 2007, Chicago, Illinois.
4. *High Performance Messaging over High-speed Networks.* University of Chicago (over Access Grid), Nov 2nd, 2006, Argonne, Illinois.
5. *SDP and Extended Sockets: A TOE and iWARP Perspective.* NetEffect Inc., June 19th, 2006, Austin, Texas.
6. *TCP and iWARP Offload Engines: Challenges and Pitfalls.* NetEffect Inc., June 19th, 2006, Austin, Texas.
7. *Designing High-performance Communication Support for Sockets-based Applications.* Hewlett-Packard (HP), May 23rd, 2006, Palo Alto, California.
8. *Enhancing the Performance and Compatibility of Sockets over High-Speed Networks.* QLogic Corporation (formerly Pathscale Incorporation), May 5th, 2006, Mountain View, California.
9. *Designing High Performance Communication Support for Sockets-based Applications over System-area and Wide-area Networks.* Argonne National Laboratory (ANL), Mar 24th, 2006, Argonne, Illinois.

Tutorials

10. *InfiniBand and 10-Gigabit Ethernet Architectures for Emerging HPC Clusters and Enterprise Datacenters.* Half-day tutorial, jointly with D. K. Panda. IEEE International Symposium on High Performance Computer Architecture (HPCA), Feb. 14, 2009, Raleigh, North Carolina.
11. *InfiniBand and 10-Gigabit Ethernet for Dummies.* Half-day tutorial, jointly with D. K. Panda and M. Koop. IEEE/ACM International Conference for High Performance Computing, Networking, Storage and Analysis (SC), Nov. 16, 2008, Austin, Texas.
12. *Designing High-End Computing Systems with InfiniBand and 10-Gigabit Ethernet.* Half-day tutorial, jointly with D. K. Panda and M. Koop. IEEE/ACM International Conference for High Performance Computing, Networking, Storage and Analysis (SC), Nov. 16, 2008, Austin, Texas.
13. *Designing HPC Clusters and Enterprise Datacenters: The InfiniBand and 10GE Way.* Half-day tutorial, jointly with D. K. Panda. IEEE Symposium on High-Performance Interconnects (HotI), Aug. 26, 2008, Palo Alto, California.
14. *High-Speed Network Architectures for Clusters: Designs and Trends.* Half-day tutorial, jointly with D. K. Panda. International Symposium on High-Performance Computer Architecture (HPCA), Feb. 16, 2008, Salt Lake City, Utah.
15. *Designing High-End Computing Systems with InfiniBand and iWARP Standards.* Full-day tutorial, jointly with D. K. Panda and S. Sur. IEEE/ACM International Conference for High Performance Computing, Networking, Storage and Analysis (SC), Nov. 10, 2007, Reno, Nevada.
16. *Designing High-End Computing Systems with InfiniBand and 10-Gigabit Ethernet.* Half-day tutorial, jointly with D. K. Panda. IEEE International Conference on Cluster Computing (Cluster), Sep. 17, 2007, Austin, Texas.
17. *Designing Clusters and Distributed Grid Computing Systems with InfiniBand and iWARP.* Half-day tutorial, jointly with D. K. Panda. IEEE International Symposium on Cluster Computing and the Grid (CCGrid), May 14, 2007, Rio de Janeiro, Brazil.
18. *State of InfiniBand in Designing HPC Clusters, Storage/File Systems, and Datacenters.* Full-day tutorial, jointly with D. K. Panda and S. Sur. IEEE/ACM International Conference for High Performance Computing, Networking, Storage and Analysis (SC), Nov. 13, 2005, Seattle, Washington.

Birds-of-a-Feather Presentations

19. *MPICH2 on DCMF*. Deep Computing Messaging Framework BoF. IEEE/ACM International Conference for High Performance Computing, Networking, Storage and Analysis (SC), 18th Nov, 2008, Austin, Texas.

Conference and Workshop Presentations

20. *Toward Message Passing for a Million Processes: Characterizing MPI on a Massive Scale Blue Gene/P*. International Supercomputing Conference (ISC), Jun. 23, 2009, Hamburg, Germany.
21. *ProOnE: A General Purpose Protocol Onload Engine for Multi- and Many-Core Architectures*. International Supercomputing Conference (ISC), Jun. 23, 2009, Hamburg, Germany.
22. *Natively Supporting True One-sided Communication in MPI on Multi-core Systems with InfiniBand*. IEEE International Symposium on Cluster Computing and the Grid (CCGrid), May 18-21, 2009, Shanghai, China.
23. *Making a Case for Proactive Flow Control in Optical Circuit-Switched Networks*. IEEE/ACM International Conference on High Performance Computing (HiPC), Dec. 20th, 2008, Bangalore, India.
24. *Sockets Direct Protocol for Hybrid Network Stacks: A Case Study with iWARP over 10G Ethernet*. IEEE/ACM International Conference on High Performance Computing (HiPC), Dec. 20th, 2008, Bangalore, India.
25. *Communication Analysis of Parallel 3D FFT for Flat Cartesian Meshes on Large Blue Gene Systems*. IEEE/ACM International Conference on High Performance Computing (HiPC), Dec. 19th, 2008, Bangalore, India.
26. *Non-Data-Communication Overheads in MPI: Analysis on Blue Gene/P*. The Euro PVM/MPI Users' Group Conference (EuroPVM/MPI), Sep. 8th, 2008, Dublin, Ireland.
27. *Impact of Network Sharing in Multi-core Architectures*. IEEE International Conference on Computer Communication and Networks (ICCCN), Aug. 4th, 2008, St. Thomas, U.S. Virgin Islands.
28. *Semantics-based Distributed I/O with the ParaMEDIC Framework*. ACM/IEEE International Symposium on High Performance Distributed Computing (HPDC), Jun. 27th, 2008, Boston, Massachusetts.
29. *Distributed I/O with ParaMEDIC: Experiences with a Worldwide Supercomputer*. International Supercomputing Conference (ISC), Jun. 17th, 2008, Dresden, Germany.
30. *Analyzing the Impact of Supporting Out-of-Order Communication on In-order Performance with iWARP*. IEEE/ACM International Conference for High Performance Computing, Networking, Storage and Analysis (SC), Nov. 14th, 2007, Reno, Nevada.
31. *Advanced Flow-control Mechanisms for the Sockets Direct Protocol over InfiniBand*. IEEE International Conference on Parallel Processing (ICPP), Sep. 13th, 2007, Xi'an, China.
32. *Analyzing and Minimizing the Impact of Opportunity Cost in QoS-aware Job Scheduling*. IEEE International Conference on Parallel Processing (ICPP), Sep. 12th, 2007, Xi'an, China.
33. *Designing NFS With RDMA For Security, Performance and Scalability*. IEEE International Conference on Parallel Processing (ICPP), Sep. 12th, 2007, Xi'an, China.
34. *High Performance MPI over iWARP: Early Experiences*. IEEE International Conference on Parallel Processing (ICPP), Sep. 12th, 2007, Xi'an, China.
35. *An Analysis of 10-Gigabit Ethernet Protocol Stacks in Multicore Environments*. IEEE International Symposium on High-Performance Interconnects (HotI), Aug. 23rd, 2007, Palo Alto, California.
36. *Nonuniformly Communicating Noncontiguous Data: A Case Study with PETSc and MPI*. IEEE International Parallel and Distributed Processing Symposium (IPDPS), Mar. 27th, 2007, Long Beach, California.

37. *Asynchronous Zero-copy Communication for Synchronous Sockets in the Sockets Direct Protocol (SDP) over InfiniBand*. Workshop on Communication Architecture for Clusters (CAC); held in conjunction with the IEEE International Parallel and Distributed Processing Symposium (IPDPS), Apr. 25th, 2006, Rhodes Island, Greece.
38. *Benefits of High Speed Interconnects to Cluster File Systems: A Case Study with Lustre*. Workshop on Communication Architecture for Clusters (CAC); held in conjunction with the IEEE International Parallel and Distributed Processing Symposium (IPDPS), Apr. 25th, 2006, Rhodes Island, Greece.
39. *Making a Case for a Green500 List*. Workshop on High-Performance, Power Aware Computing (HP-PAC); held in conjunction with the IEEE International Parallel and Distributed Processing Symposium (IPDPS), Apr. 25th, 2006, Rhodes Island, Greece.
40. *Head-to-TOE Comparison for High Performance Sockets over Protocol Offload Engines*. IEEE International Conference on Cluster Computing (Cluster), Sep. 29th, 2005, Boston, Massachusetts.
41. *Supporting RDMA Capable Network Compatibility and Features for Regular Network Adapters*. Workshop on Remote Direct Memory Access (RDMA): Applications, Implementations and Techniques (RAIT); held in conjunction with the IEEE International conference on Cluster Computing (Cluster), Sep. 26th, 2005, Boston, Massachusetts.
42. *Performance Characterization of a 10-Gigabit Ethernet TOE*. IEEE International Symposium on High Performance Interconnects (HotI), Aug. 18th, 2005, Palo Alto, California.
43. *Sockets vs. RDMA Interface over 10-Gigabit Networks: An In depth Analysis of the Memory Traffic Bottleneck*. Workshop on Remote Direct Memory Access (RDMA): Applications, Implementations and Technologies (RAIT); held in conjunction with the IEEE International Conference on Cluster Computing (Cluster), Sep. 20th, 2004, San Diego, California.
44. *Exploiting Remote Memory Operations to Design Efficient Reconfiguration for Shared Data-Centers over InfiniBand*. Workshop on Remote Direct Memory Access (RDMA): Applications, Implementations and Technologies (RAIT); held in conjunction with the IEEE International Conference on Cluster Computing (Cluster), Sep. 20th, 2004, San Diego, California.
45. *Towards Provision of Quality of Service Guarantees in Job Scheduling*. IEEE International Conference on Cluster Computing (Cluster), Sep. 22nd, 2004, San Diego, California.
46. *Sockets Direct Protocol over InfiniBand in Clusters: Is it Beneficial?*. IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS), Mar. 10th, 2004, Austin, Texas.
47. *Impact of High Performance Sockets on Data Intensive Applications*. IEEE International Symposium on High Performance Distributed Computing (HPDC), Jun. 22nd, 2003, Seattle, Washington.
48. *Efficient Collective Operations using Remote Memory Operations on VIA-based Clusters*. IEEE International Parallel and Distributed Processing Symposium (IPDPS), Apr. 24th, 2003, Nice, France.
49. *High Performance User-level Sockets over Gigabit Ethernet*. IEEE International Conference on Cluster Computing, Sep. 25th, 2002, Chicago, Illinois.

Instructional Seminars

50. *Message Passing for Dummies: Introduction to MPI*. Student Lecture Series, Argonne National Laboratory, Jun 9th, 2008, Argonne, Illinois.

Demo/Project Showcase

51. *ParaMEDIC: Parallel Metadata Environment for Distributed I/O and Computing*. IEEE/ACM International Conference for High Performance Computing, Networking, Storage and Analysis (SC), Nov 13th to 15th, 2007, Reno, Nevada.
52. *High-performance Message Passing with MPICH2*. Argonne National Laboratory exhibition booth at the IEEE International conference on Supercomputing (SC), Nov 14th to 16th, 2006, Tampa, Florida.

53. *mpiBLAST on the GreenGene Distributed Supercomputer: Sequencing the NT database against the NT database (An NT-complete problem)*. In conjunction with the Los Alamos National Laboratory at the IEEE International conference on Supercomputing (SC), Nov 14th to 17th, 2005, Seattle, Washington.
54. *InfiniBand: Performance and Implications on Next Generation Applications*. Exhibition Demo at the 10th Annual Coalition for National Science Funding Science Exhibition and Reception, Jun 23rd, 2003, Capitol Hill, Washington DC.

Poster Presentations

55. *Semantics-based Distributed I/O for mpiBLAST*. ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP), Feb. 20-23, 2008, Salt Lake City, Utah.

Miscellaneous Presentations and Seminars

56. *Analyzing, Optimizing and Evaluating PETSc over MPI*. IEEE International conference on Supercomputing (SC), Argonne National Laboratory Exhibition Booth, Nov 14th, 2006, Tampa, Florida.
57. *High-performance Communication Support for Sockets-based Applications*. Ohio State University (OSU), June 15th, 2006, Columbus, Ohio.
58. *Comparing 10-Gigabit Ethernet with InfiniBand and Myrinet for High Performance Sockets over Protocol Offload Engines*. IEEE International conference on Supercomputing (SC), ASC Exhibition Booth, Nov 16th, 2005, Seattle, Washington.
59. *A Case for UDP Offload Engines in Lambda Grids*. Ohio State University (OSU), Nov 4th, 2005, Columbus, Ohio.
60. *Analyzing the Memory Traffic Bottleneck in TCP/IP over 10-Gigabit Ethernet*. Ohio State University (OSU), Sep 12th, 2002, Columbus, Ohio.
61. *Software iWARP: Design and Implementation in FreeBSD*. Ohio State University (OSU), Sep 13th, 2002, Columbus, Ohio.
62. *Java Animated Text Rendering for Indian Language Scripts*. Indian Institute of Technology (IIT), May 11th, 2001, Madras, India.

PROFESSIONAL ACTIVITIES

Chairmanships and Editorships

1. Journal Special Issue Co-editor (with Dr. Abhinav Vishnu): *Parallel Programming Models and Systems Software for High-End Computing Applications*, Special Issue with the International Journal of High Performance Computing Applications (IJHPCA), 2009.
2. Program Co-chair (with Dr. Ada Gavrilovska): *International Workshop on High Performance Interconnects for Distributed Computing (HPI-DC)*, 2009.
3. Tutorials Chair: *IEEE International Conference on High Performance Interconnects (HotI)*, 2009.
4. Journal Special Issue Co-editor (with Dr. Wu-chun Feng): *Tools and Environments for Multi- and Many-Core Architectures*, Special Issue with IEEE Computer, 2009.
5. Program Co-chair (with Dr. Abhinav Vishnu): *International Workshop on Parallel Programming Models and Systems Software for High-end Computing (P2S2)*, 2009.
6. Track Co-chair (with Dr. Li Xiao): *IEEE International Conference on Computer Communications and Networks, Pervasive Computing and Grid Networking Track*, 2009.
7. Program Co-chair (with Dr. Sayantan Sur): *International Workshop on Parallel Programming Models and Systems Software for High-end Computing (P2S2)*, 2008.

Technical Committees

1. *Program Committee for International Conferences and Workshops:*

- (a) International Workshop on The Influence of I/O on Microprocessor Architecture (IOM): 2009.
- (b) International Workshop on Communication Architecture for Clusters (CAC): 2009.
- (c) IEEE International Symposium on High-Performance Interconnects (HotI): 2008, 2009.
- (d) International Conference on Parallel Processing (ICPP): 2008, 2009.
- (e) International Workshop on Scheduling and Resource Management for Parallel and Distributed Systems (SRM-PDS): 2007, 2008, 2009.
- (f) IEEE International Conference on High Performance Computing (HiPC): 2007, 2008, 2009.
- (g) IEEE International Conference on Computer Communications and Networks (ICCCN): 2006, 2007, 2008, 2009.

2. *Posters Committee:*

- (a) IEEE International Conference on Cluster Computing (Cluster): 2009.

3. *Tutorial Committee:*

- (a) IEEE/ACM International Conference for High Performance Computing, Networking, Storage and Analysis (SC): 2009.

4. *Technical Review Committee for International Journals and Magazines:*

- (a) Journal of Concurrency and Computation: Practice and Experience: 2009.
- (b) Computer Networks Journal: 2009.
- (c) IBM Systems Journal: 2008.
- (d) IEEE Micro: 2007.
- (e) IEEE Transactions on Parallel and Distributed Systems (TPDS): 2006, 2007, 2008, 2009.
- (f) IEEE Transactions on Computers (TC): 2006, 2007, 2008.
- (g) Electronics and Telecommunications Research Institute Journal (ETRI): 2006.
- (h) Journal of Grid Computing (Grid): 2005.
- (i) Journal of Parallel and Distributed Computing (JPDC): 2004, 2005, 2006, 2007, 2008, 2009.

5. *Organizational Positions:*

- (a) Session Chair: IEEE International Conference on High Performance Computing (HiPC): 2008.
- (b) Publicity Chair: Workshop on High-Performance, Power-Aware Computing (HP-PAC): 2008.
- (c) Session Chair: IEEE/ACM International Conference for High Performance Computing, Networking, Storage and Analysis (SC): 2007.
- (d) Session Chair: International Conference on Parallel Processing (ICPP): 2007.
- (e) Session Chair: Workshop on Performance Optimization for High-Level Languages and Libraries (POHLL): 2007.
- (f) Session Chair: Workshop on Communication Architecture for Clusters (CAC): 2007, 2008.
- (g) Session Chair: IEEE International Conference on Computer Communications and Networks (ICCCN): 2006, 2008.

6. *Additional Reviewer for Conferences and Workshops:*

- (a) IEEE International Conference on Computer Communications (INFOCOM): 2009.
- (b) IEEE International Symposium on High-Performance Interconnects (HotI): 2007.

- (c) ACM International Conference on Supercomputing (ICS): 2007.
- (d) IEEE International Conference on Autonomic and Trusted Computing (ATC): 2007.
- (e) IEEE/ACM International Conference for High Performance Computing, Networking, Storage and Analysis (SC): 2005, 2007, 2008.
- (f) IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS): 2005.
- (g) IEEE Symposium on High-Performance Computer Architecture (HPCA): 2005.
- (h) IEEE Symposium on High Performance Distributed Computing (HPDC): 2004, 2005, 2007, 2008.
- (i) IEEE International Parallel and Distributed Processing Symposium (IPDPS): 2004, 2005, 2006.
- (j) International Workshop on RDMA: Applications, Implementations and Technologies (RAIT): 2004, 2005.
- (k) IEEE International Conference on Cluster Computing: 2003, 2004, 2005.
- (l) International Workshop on Communication Architecture for Clusters (CAC): 2003, 2004, 2005.

Panel Moderator

1. *Software for the ExaFlop Era: The Demons & The Dementors*. First International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2), 2008.

Advisees (Interns & Postdocs)

1. James Dinan. *Analyzing Hybrid MPI-UPC Programming*. 4th year Ph.D. student, Dept. of Computer Science and Engineering, Ohio State University, 2009.
2. Gopalakrishnan Santhanaraman. *Truly One-sided RMA Operations in MPI*. 7th year Ph.D. student, Dept. of Computer Science and Engineering, Ohio State University, 2008.
3. Ping Lai. *ProOnE: A General Purpose Protocol Onload Engine*. 1st year Ph.D. student, Dept. of Computer Science and Engineering, Ohio State University, 2008.
4. Rajesh Sudarsan. *A Component-based Process Manager for MPICH2*. 4th year Ph.D. student, Dept. of Computer Science, Virginia Tech., 2008.
5. Thomas Scogland. *Efficient Boot-strap Capabilities for MPI Processes*. 1st year Ph.D. student, Dept. of Computer Science, Virginia Tech., 2008.
6. Ganesh Narayanaswamy. *MPICH2 over Qlogic PSM*. 1st year M.S. student, Dept. of Computer Science, Virginia Tech., 2007.

Thesis Committees

1. Ganesh Narayanaswamy. *On the Interaction of High-Performance Network Protocol Stacks with Multicore Architectures*, Virginia Tech., M.S., 2008.

Other Committees

1. Student Research Symposium Poster Awards Committee at the IEEE International Conference on High Performance Computing (HiPC), 2008.
2. Book Proposal Review Committee, Taylor and Francis Group, 2007.
3. Departmental Awards Committee, Computer Science and Engineering, Ohio State University, 2006.

Professional Memberships

1. Member, IEEE Computer Society
2. Member, ACM Computer Society

REFERENCES

Available on Request